Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story
Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles
Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor
To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives
Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking
We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

Risks and Opportunities in ASIC Design

Craig Peterson
Co-General Manager
Intel Microelectronics Services
Intel Corporation

Overview

The use of standard silicon building blocks is increasing in the networking and communications industry. Standard parts help to reduce product development and manufacturing costs and to improve interoperability across vendors and technologies. Of course, companies still need to differentiate their products, and one method is to integrate a customized ASIC to support specialized functions.

In creating these ASICs, the demanding requirements of networking and communication applications typically require the use of 0.18- or 0.13-micron process technologies. Yet ASIC design practices that were successful on earlier process generations are not sufficient to deal with these complex, high-density solutions. Advanced tools and methodologies are needed to address key challenges that have emerged as by-products of complex designs and high-density scaling.

This article describes the scaling challenges and offers suggestions for evaluating the ability of an ASIC vendor to deliver complex, high-density ASICs on spec, on time, and within budget. For more detailed information, see the related Intel white paper, or visit the Intel® Microelectronics Services Web site.

The Seven Key Challenges of ASIC Design

The microelectronics industry is facing a major and unprecedented shift in ASIC design requirements. Today's advanced process generations have achieved densities and complexities at which new effects come into play. As a result, standard methodologies can no longer deliver reliable designs on a predictable schedule. Advanced tools and methods have been developed to deal with these problems, yet many are not yet widely used in the ASIC industry.

Altogether, there are seven key design challenges that have arisen due to the effects of scaling. Each challenge is discussed in a section that follows.

The Process Challenge

A variety of physical side effects began to emerge with the 0.25-micron process generation, and have become more challenging with each successive generation. A successful design flow must incorporate advanced design tools to address each of these issues:

- **RC Coupling** *(0.25 micron and smaller).* Interconnect resistance and capacitance parasitics can account for up to 80 percent of the delay on long top-level interconnect signals, resulting in thousands of potential timing violations.
- **Signal Integrity** *(0.18 micron and smaller).* Cross capacitance between signals can cause a logic error in a tight speed path. Even if early manufacturing runs are fully functional, signal integrity problems can arise with even a small amount of manufacturing process drift in later runs.
- **In-die Variation** *(0.13 micron and smaller).* Variations in transistor strength across each die can make it difficult to support a common clock domain across high-frequency chips without proper margining and block-to-block timing. Logical clock domains often have to be divided into multiple physical clock domains that are linked with well-controlled timing interfaces.
- **Transistor Leakage** *(0.13 micron and smaller).* As geometries and threshold voltages are scaled down, transistor leakage can become a recognizable percentage of the total power budget of a chip. Improved design techniques can reduce overall leakage power.
- **Soft Error Rate** *(0.13 micron and smaller).* The energy stored in static memory nodes has become small enough that a stray alpha particle can cause a latch to flip states. Advanced design tools can detect potential problems and provide alternatives.

The Complexity Challenge

Today's most complex ASIC designs can have 5M or more gates. Yet for a typical design with only 800K gates, it may take a week to run a full back-end iteration using traditional, flat design methods. As a result, even a few unplanned iterations can add several weeks to the design cycle, making schedules unpredictable.

New hierarchical design techniques partition a design into smaller physical blocks, which can be run independently and in parallel. With this approach, a design iteration can be performed in about a day. Schedules are faster and more predictable. Of course, advanced tools, methodologies, and expertise are required to take advantage of hierarchical design strategies.

The Performance Challenge

High-speed functions used to be supported on separate chips. Now they are being integrated on the ASIC to improve overall system performance and reduce product costs. Examples include high-speed serial interfaces such as Ethernet, SONET, and FibreChannel.

These high-speed functions raise a variety of new issues when clock frequencies go beyond 300–400 MHz. Hold time becomes significantly more important due to in-die variation and the decreased number of gates between flip-flops. Impedance, signal cross talk, and RC coupling become increasingly serious concerns. Flip-chip die attach may be required to reduce die-package interconnect inductance. Decoupling at die, package, or board should be designed to eliminate power supply oscillation, and considerable care must be taken in routing high-frequency signals. Routing delays can also become problematic, and must be included, along with gate delays, for accurate timing analysis.

Interconnect performance can be improved significantly by:

- Pre-routing critical signals so they have preferred routes
- Inserting repeaters over long signal paths
- Making efficient use of the additional metal layers that are available in advanced process generations

Signal integrity analysis should also be included in early design iterations. New tools and methodologies are required to integrate these options smoothly into the design process.

The Power Challenge

Power supply voltages are dropping about 30 percent with each new process generation. Yet power consumption is doubling with each generation due to increases in frequency and logic density. Leakage power is also increasing because of shrinking transistor threshold voltages.

Falling power supply voltages and rising power consumption is causing dramatic increases in power supply current. The total power supply current for high-end designs may now exceed 22 amperes. With these currents, even a little resistance in the voltage distribution system can significantly impact performance.

Since power issues are not addressed in a normal timing analysis roll-up, new methodologies are needed to accommodate the effects of leakage power and to ensure power supply integrity. Each of the following options can be used, but they must be efficiently integrated into design tools and processes.

- **Process selection.** Foundries now offer multiple options on the 0.13-micron process generation, such as normal, low power, low voltage, high V_t, and low V_t.
- **Cell library selection.** Every library design involves a variety of unique design trade-offs, so the choice of library can be critical.
- **Clock Tree Design.** Clock gating can be incorporated to turn off portions of the clock tree for inactive logical blocks.
- **Logic Synthesis.** Most current design methodologies use standard wireload models that can result in oversized local logic that increases power consumption, or in undersized drivers for longer signals. Advanced tools can optimize synthesis, sizing, and initial cell placement to address these issues.
- **Power Grid.** Power delivery can be improved by using upper layers to form a power grid that can tie directly to cell row power (common methodologies use power rings around blocks).
- **Flip Chip.** For designs with extreme core power requirements, lots of I/O, or very high-performance I/O, flip-chip package assembly can provide better signaling characteristics and improved power distribution to the die.

The Density Challenge

Layout assembly by abutment—along with the use of additional metal layers in advanced process generations—can significantly improve functional density in ASIC designs. Other factors in optimizing density include:

- **Cell Library Selection.** Cell row height is a key parameter of cell library design, and can have a significant impact on density. Additional factors complicate the issue, requiring sophisticated analysis of library options. A library that works well on one design may not have the characteristics that make it work well for the next design.
- **Optimum Synthesis, Sizing, and Placement.** Density can be improved by optimizing the size and placement of cells in blocks to minimize wire resources and delays. Advanced tools can improve density as well as performance by performing gate-level synthesis and gate placement at the same time.
- **Datapath Centric Design Optimization.** Conventional place and route tools do not differentiate between a datapath and its control. Advanced tools can be used to explore a variety of microarchitectural approaches, and to achieve placed and routed designs that rival hand-packed datapath techniques.

The Quality and Reliability Challenge

Advanced manufacturing processes create new sources of potential chip failure. One example is SER (soft error rate), logic errors caused by alpha particles or other subatomic events in static memory and flip-flops as well as in DRAM memory.

In earlier process generations, similar physical effects were negligible. Unless the design methodology can detect and resolve these issues, an increasing number of parts may pass debug tests and then fail in the field.

The Productivity Challenge

Time-to-market remains a driving force in the communications industry, so ASIC development must stay off the critical path in system development schedules. Over the past 20 years, the complexity of ASIC design has grown much faster than development capabilities. Significant progress has been made recently, but even the very latest EDA tools cannot close this “design productivity gap.” To address this issue, ASIC suppliers need to:

- Develop more efficient, hierarchical flows that partition ASICs into manageable segments that can be designed in parallel with faster turnaround.
- Re-engineer methodologies, tools, and skill sets to address the increasing number of deep submicron issues.
- Become proficient in rapidly integrating internal or commercially available IP cores.

Recommendations for Reliable ASIC Design

Because of the seven challenges outlined in this article, developers planning to integrate ASICs into their networking and communication products should look closely at the tools, methodologies, and expertise of an ASIC vendor. If the ASIC will be implemented using 0.13-micron process generation and beyond, advanced methods will be required to ensure predictable, timely delivery of a reliable product. Techniques and methods that worked previously may not be successful on the next process.

When assessing an ASIC vendor, each of the seven challenges should be considered. The following questions offer a starting point for exploring vendor capabilities for deep submicron design:

- How is the design partitioned to ensure development efficiency?
- Is block partitioning different for synthesis than for place and route?
- How is the optimal shape and size determined for each block partition?
- How are boundary signal crossing points planned for each block, so they support the shortest possible inter-block routing and the best in-block cell placement for global optimization?
- How are global repeaters placed and routed to maximize productivity and timing convergence, while providing the best possible interconnect performance?
- How are timing interfaces implemented between blocks, and how is the timing rolled up to support accurate timing analysis, keeping in mind such issues as in-die variation?
- How are top-level clock buffering and routing implemented, as opposed to local block buffering and routing?
- How is top-level power planned and interconnected to block-level power and made electrically correct?
- What kind of interconnect resource sharing can be achieved between top-level and block-level layers to reduce congestive areas, and to increase speed on critical speed paths?
- What type of pre-layout is performed to optimize critical signals?

- How are signal integrity checks and fixes performed?
- What steps are taken to ensure signal, power, and timing integrity for high-frequency functions?

Summary

Seven key challenges now confront ASIC design teams as they work to create faster and more complex ASICs using 0.18 micron, 0.13 micron, and smaller process generations. Advanced tools and methodologies are available to address these challenges, and to enable reliable ASIC design on a predictable schedule. Developers should protect themselves and their product timelines by carefully evaluating their ASIC vendor's ability to handle the level of performance and complexity required for their particular application.

More Info

- For more information on advanced tools and methodologies for ASIC design, visit the Intel Microelectronics Services Web site.
- For more information about scaling challenges, download the Intel Microelectronics Services white paper, "Seven Critical Scaling Challenges of ASIC Design."

Author Bio

Craig Peterson has worked at Intel in design and design technology for 27 years, and is currently the co-general manager for Intel Microelectronics Services. As technology director in design technology, he drove dramatic improvements in microprocessor and chipset development productivity and optimization and performance across much of Intel.

Previously, Craig led the design and development of two processors and five generations of chipsets. He also co-invented and headed the interconnect component development for the world's first supercomputer to break the sustained TeraFLOPS computation barrier.

Craig has achieved many firsts at Intel. He co-wrote the first RTL simulation, wrote the first cell-based synthesis tool, drove the first all-workstation-based project, designed the first single-chip memory controller, designed the first single-board multiprocessor board, and co-microarchitected the first multithreaded processor.

Departments

Desktop

Intel® 845 Chipset and Desktop Boards Open New Markets with DDR Support

Heath Winston
Discrete Chipset Business Manager
Desktop Platforms Group
Intel Corporation

Scott Diaz
Senior Product Marketing Engineer
Desktop Product Solutions Division
Intel Corporation

Overview

No matter what level of memory desktop integrators may want to provide, no matter what market segment they are targeting, Intel now can provide them a complete platform solution. As before, Intel continues its support for the high-performance RDRAM with the Intel® 850 chipsets and boards and its support for the high-value SDR-SDRAM with the Intel® 845 chipsets and desktop boards. Additionally, Intel now supports the midrange and increasingly popular DDR-SDRAM in its 845 chipset and provides it on its 845 boards. Consequently, integrators using Intel® platforms have the flexibility to meet a wider-than-ever range of price/performance points and to increase volume and reduce cost.

In this article you will learn about the Intel 845 chipset and its support for DDR-SDRAM, the three boards into which the 845 chipset is integrated, and the features that make each of these boards ideally suited for a given market segment.

DDR in a "Full Benefits" Package

Other chipset manufacturers have offered support for DDR-SDRAM (also known as DDR), but Intel decided to wait until the technology could be enabled as highly reliable, interoperable, and sufficiently available for a high-volume ramp. By extending its memory support to DDR, the Intel 845 chipset provides integrators a full benefits package: the ability to meet the vast mainstream market segment between purchasers of systems based on RDRAM and those based on SDR-SDRAM (also known as SDRAM) along with the performance, reliability, and support that Intel is known for.

The Intel 845 chipset is designed, validated, and optimized for the Intel® Pentium® 4 processor with Intel® NetBurst™ microarchitecture, extending the reach of that powerful platform into mainstream markets. To optimize DDR memory performance, the 845 chipset implements enhancements that minimize setup times and maximize throughput. For example, the 82845 Memory Controller Hub supports both 200/266 DDR (and 100/133 SDRAM) and the latest graphics devices through a 1.5V AGP4X interface. The 82801BA I/O Controller Hub connects graphics and memory directly for faster peripheral access. A 400-MHz system bus provides three times the bandwidth of previous system-bus technologies, and the new SSE2 instruction set ensures a robust foundation for current and emerging software, especially the high-end audio, video, and 3D gaming applications favored by consumers.

A Board for Every Market Segment

To provide the most comprehensive support for the consumer and selected business market segments requiring DDR, Intel integrates its 845 chipset in three desktop boards: the D845BG, the D845PT, and the D845WR. These boards offer a variety of form factors, networking and expansion capabilities, and peripheral support.

The Intel® Desktop Board D845BG is available in three models: the D845BG, the D845BGL, and the D845BGSE, all based on the ATX form factor and including 1.5V 4X/2X AGP, two-channel AC'97 audio, and five or six PCI slots. The D845BG and D845BGSE provide a CNR (Communications and Networking Riser), which, in combination with Intel® chipsets, provides cost-effective riser card solutions for surround-sound audio, modem, and other features ideal for consumer segments. The D845BGL provides an Intel® PRO/100 Network Connection, particularly suited for business segments requiring handy LAN capabilities.

The D845BG and D845BGL also provide, respectively, seven and six USB 1.1 ports, targeting more cost-conscious markets using standard peripherals, while the D845BGSE provides up to five USB 2.0 ports. The high-speed USB 2.0 interface provides up to 40 times the bandwidth of USB 1.1, making the D845BGSE a perfect fit for consumer markets demanding high-speed access to the latest scanners, cameras, audio devices, CD-ROM burners, and other powerful peripherals. For integrators targeting technology-savvy consumers, the D845BGSE is the ideal board for a leading-edge mainstream system.

The Intel® Desktop Board D845PT is available in two models: the D845PT and the D845PTL. These boards are identical, respectively, to the D845BG and D845BGL except that the D845PT and D845PTL are based on the microATX form factor, and therefore support just three PCI slots. Because of its very compact size, the D845PT is ideally suited to business markets that want the performance and reliability of the Intel Pentium 4 processor and Intel 845 chipset but do not need the expandability of an ATX board. The D845PT is also well suited for a range of consumer systems, from low-cost systems to innovative small form factor designs.

The Intel® Desktop Board D845WR provides some different features from the D845PT or D845BG to address the needs of other market segments. Features such as a MIDI game port, an additional serial port on the back panel, and an ISA slot, make the board well suited for China and other markets desiring these I/O options. The D845WR is also available with integrated six-channel AC'97 audio providing a low-cost surround-sound solution for retail and other consumer channels.

In addition to the features specific to the D845BG, D845PT, and D845WR, respectively, all the boards provide wider data paths and flexible memory-refresh technology for the optimum DDR SDRAM performance; advanced packaging technology and innovative electrical design to ensure long-term reliability over a variety of operating conditions; dual UltraATA/100 controllers for faster IDE transfers to and from storage devices; Intel® Rapid BIOS Boot, for faster system access; and Instantly Available PC, Intel® Active Monitor, Alert on LAN, and low-power sleep mode, for cost-effective and energy-efficient operation.

The Intel Desktop Boards D845BG, D845PT, and D845WR also come with a comprehensive package of value-add software designed to meet the demands of small businesses launching e-Commerce projects, home users creating personal music libraries, and just about any application in between. These software products include Norton Internet Security* (security and virus protection), NTI CD-Maker 2000* (full-featured CD-ROM creation), SoundMAX* with SPX (high-performance audio), RealPlayer* (streaming audio and video), RealJukebox* (recording and listening to MP3 audio), Power Technology DFX* (enhancing the performance of RealPlayer and RealJukebox), and MacroMedia Shockwave* (3D Internet).

Summary

Now that DDR has matured to become a memory technology as stable and available as its predecessors, SDRAM and RDRAM, Intel supports it enthusiastically in the Intel 845 chipset and provides it on the boards integrating that chipset. This means that system developers and integrators now have a full package based on the Pentium 4 architecture—processor, chipset, desktop board, and value-add software—supporting the memory demands of any market segment they might wish to target. It also means they can enjoy this enhanced flexibility along with the traditional advantages of the Intel environment: extensive validation and testing, reliability, availability, and three-year limited warranty for Intel® motherboards.

More Info

For more information on the Intel 845 chipset and Desktop Boards D845BG, D845PT, and D845WR, visit the Intel® Chipsets area or the Intel® Desktop Boards area of the Intel Developer Site.

Author Bios

Heath Winston is a discrete chipset business manager in the Chipset Business Management Marketing division of the Desktop Platforms Group at Intel Corporation, where he has worked for five years. Most recently, he has focused his efforts on the launch of the Intel 845 and 850 chipsets. Heath holds a B.S. in business administration from the California State University at Sacramento.

Scott Diaz is a senior product marketing engineer in the Desktop Product Solutions Division at Intel Corporation, where he has worked for two years. Most recently, he has focused his efforts on the launch of the Intel 845 chipset and Intel® 815E products and has led roadmap planning cycles and development of the Board Data Center. Scott holds an M.B.A. from the University of Chicago and a B.S. in mechanical engineering from Iowa State University.

IDE RAID on the Desktop

Scott Jackson
Technical Marketing Manager
Desktop Platform Solutions Division
Intel Corporation

Overview

Since 1997, CPU clock frequencies have increased tenfold and memory bandwidth sixfold, while hard-drive performance has grown a mere threefold. The effect of this trend is best explained in Patterson, Gibson, and Katz's opening of their 1988 "A Case for Redundant Arrays of Inexpensive Disks (RAID)" paper, introducing RAID. "Increasing performance of CPUs and memories will be squandered," they stated, "if not matched by similar performance increase in I/O." Their paper proposed several levels of RAID, all aimed at increasing disk I/O performance and reliability.

Until recently, the performance and reliability of RAID were available only to a limited few willing to pay the high price carried by SCSI-based systems. The introduction of IDE RAID–Integrated Development Environment RAID–has allowed for a cost-effective RAID solution, broadening market reach into the mainstream, low-end server, and workstation market segments.

Striping and Mirroring: the Advantage

RAID is an acronym for Redundant Array of Inexpensive Drives, first proposed by Patterson, Gibson, and Katz of the University of California, Berkeley. RAID is available in a range of functionalities addressing diverse performance and fault-tolerance requirements, but by far the most common in the desktop and low-end server arena are RAID level 0 (or, simply, RAID 0) and RAID level 1 (or, simply, RAID 1).

There are several reasons why these are the most common levels of RAID. First, these levels are basic and require very little overhead; therefore, additional hardware is not necessary, keeping the cost of implementations down. Second, more advanced RAID levels, requiring additional hardware, are basically variants based on the principles of data redundancy (RAID 1), enhanced performance (RAID 0), or both. Therefore higher RAID level functionalities can be somewhat replicated with these basic levels.

RAID 0, also known as striping, uses two or more physical drives to create a single logical drive. It works by distributing data in blocks assigned to given drives and reading the data from or writing it to each of these drives concurrently. For applications relying on large files, RAID 0 significantly boosts performance by increasing the available I/O bandwidth, a limiting factor in large file transfers.

Theoretical bandwidth is increased by a factor of *n*, where *n* is the number of drives used. Real-world performance of desktop systems incorporating IDE RAID 0 with two hard drives showed an overall 37 percent performance improvement as measured by WinBench99* High-End DiskMark* and an overall 34 percent performance improvement in the Business DiskMark. Common applications used on high-end desktops and low-end workstations such as the Adobe* Premiere* 4.2 video-editing software showed an overall 32 percent performance improvement, and Adobe PhotoShop* 4.0 showed an overall 34 percent performance improvement.[1]

The tradeoff with RAID 0 is fault tolerance. Because it distributes data across multiple hard drives, a failure in any one of the drives will result in an effective loss of all the data. This combination of stronger performance but weaker fault tolerance makes RAID 0 well suited for high-data-volume applications such as video, music, or photo editing, but poorly suited for mission-critical applications such as an enterprise database. For such environments and applications, RAID 1 is a better solution.

Like RAID 0, RAID 1, also known as mirroring, uses two or more drives. But unlike RAID 0, RAID 1 uses these drives for data duplication instead of data distribution. For any given set of data, RAID 1 writes to both (all) drives concurrently, effectively storing *n* copies of the data, where *n* is the number of drives used. RAID 1 is an ideal solution for users and businesses running mission-critical applications and needing an easily implemented and cost-effective approach to fault tolerance.

The tradeoff with RAID 1 is drive capacity. Because the data on each drive in a RAID 1 array is a duplicate, the hard-drive capacity is equal to the smallest capacity drive in the RAID array. For example, in a RAID 1 setup with a 20-GB drive and a 30-GB drive, the disk subsystem would appear to the user as a single 20-Gb hard drive. RAID 1 also offers little to no performance increases over a single-drive setup.

Common RAID Implementations

Already, a number of major computer manufacturers and top integrators are incorporating RAID into desktop and low-end server systems. Following are common implementations of RAID:

Software RAID is implemented with an additional IDE controller on the motherboard or a PCI card. Because it requires an additional hardware component, Software RAID is sometimes confused with Hardware RAID. However, it differs significantly from Hardware RAID in that overhead is processed by the CPU instead of a separate microcontroller. This is the most common implementation of IDE RAID among major computer manufacturers, top integrators, and enthusiasts. The majority of systems shipping with IDE RAID use this implementation.

Hardware RAID is implemented with a single or multiple IDE controllers, a microcontroller, and dedicated memory on the motherboard or a PCI card. Hardware RAID is less common than Software RAID on desktops and low-end servers because it is more costly and generally does not offer performance levels equivalent to those of Software RAID[2]. What Hardware RAID does offer is additional levels of RAID. These additional levels require dedicated processing power supplied by the integrated microcontroller. An example would be RAID level 5, which offers the performance of striping and the fault tolerance of mirroring. Although this level could be run without additional hardware, to realize its full potential and performance a dedicated microcontroller is required.

OS Software RAID is implemented through the existing I/O controller on the motherboard by the Windows NT*, Windows* 2000, or Windows XP operating system. The advantage of OS RAID is that additional hardware is not required; the two IDE channels supplied with motherboards can be used. Initially this looks to be the least expensive implementation but there are many drawbacks. First, each version of Windows, from Windows NT Server to Windows XP Professional, offers different levels of RAID support. While Windows NT Server supports RAID 0, 1, and 5, NT Server's price premium dwarfs the cost of a software RAID setup. Windows XP Professional supports only RAID 0 while the home edition doesn't support RAID at all. Additionally, to take advantage of the performance offered in RAID 0, each hard drive should have a dedicated IDE channel. This ensures there won't be peripherals competing with the hard drive for control of the IDE channel, potentially slowing data transfers down. The price premium of a Windows OS with RAID support coupled with the possibility of needing an additional IDE controller to fully realize RAID's potential makes this a less than ideal implementation for the desktop market segment.

Cost

Integrators wishing to add Software RAID to their system offerings have a couple choices to make. First, will you add the additional IDE controller integrated into the motherboard or offer an optional PCI RAID controller card? An integrator can assume a $10 to $20 (USD) price premium over identical motherboards without RAID, while a PCI card can range from $60 to $75 (USD). The next cost associated with RAID systems is an additional hard drive. The cost of two hard drives can actually save an integrator money depending on the size of the hard drives. For example, an IBM* 60-Gb 7,200 RPM drive can be purchased for $111 (USD) online; an identical drive from the same manufacturer but 120 Gb in size is $288 (USD)[3]. In this situation an integrator can offer the same size hard drive, 120 Gb, with the increased performance of RAID 0, for less than what a single drive would cost. The above case demonstrates that the performance advantage of RAID 0 can potentially be had for little to no cost to system integrators.

Summary

The promise of increased speed and reliability at little to no cost is sure to push RAID deep into the mainstream market segments. Introduction of cost-effective IDE RAID has businesses running mission-critical applications on systems that cost fractions of what they used to. Photo and video editing once possible only on ultra-high-end workstations running SCSI RAID setups can now be done on entry-level workstations and high-end desktops. IDE RAID has enabled new markets for integrators and has brought performance and reliably to a new price point.

References

1. System configuration–Source: Intel Corporation
Intel® Pentium® 4 Processor @ 2.0 GHz (256-Kb Cache), Intel® Desktop Board D850MD, BIOS MV580.10A.86A.0008.P04, Chipset INF v3.20.1008, 2x 256 MB PC800 RDRAM*, ATA-100, 2 x Seagate* Barracuda* ATA III Model ST320414A 2-MB cache, Promise* Fasttrack* 100TX2 (PDC20270 controller), VisionTek* GeForce3* 64-MB video card, Detonator* 21.81 video drivers, 1,024x768x16 bpp–75 MHz, Microsoft* DirectX* 8.0, On-board LAN disabled, IDE Storage drivers: Intel® 82801BA ICH2 6.20.2018, Promise drivers Ver2.0, Windows* 2000 Service Pack 1, FAT32 file system.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel® products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference this Benchmark Limitations page.

Benchmark Application: eTesting Labs Inc.'s WinBench* 99 Version 2.0
Test was performed without independent verification by eTesting Labs Inc. and that eTesting Labs Inc. makes no representations or warranties as to the result of the test.

Note: The WinBench99 benchmark is the latest WinBench offering from eTesting Labs Inc. The advisory committee of ZDNet editors decided not to offer a 2001 version of WinBench, but rather update WinBench99 to version 2.0.

2. “IDE RAID Comparison”–AnandTech.com

3. http://www.newegg.com –1/17/02

- IBM EIDE HARD DRIVE 120-GB 7,200 RPM MODEL # IC35L120AVVA07 and IBM EIDE HARD DRIVE 60-GB 7,200 RPM MODEL # IC35L060AVER07
- IWILL* SIDE-RAID 100 High Point HPT370 I/O CONTROLLER / PROMISE FASTTRAK*100 TX2 RAID CONTROLLER CARD
- ABIT* TH7II INTEL® 850 ATX MOTHERBOARD / ABIT TH7II-RAID INTEL 850 ATX MOTHERBOARD

Author Bio

Scott Jackson is currently a member of the Intel Sales and Marketing Rotation Program. His current assignment is with the Desktop Platform Solutions Division as a Product Technical Marketing Engineer. He also has held product marketing engineer positions in the same division and in the Cellular Communication Division. Before joining Intel in July 2000, Scott worked for Digital Equipment Corporation and Compaq Computer Corporation as a hardware engineer on the Alpha processor family. He holds a B.S.E. in computer engineering from the University of Michigan.

Initiatives & Technologies

Intel Press Unveils Five New Technical Books

Richard Bowles
Publisher and Managing Director
Intel Press
Intel Corporation

Overview

Intel Press has announced five new technical books aimed at streamlining the development and deployment of systems and software applications based on the latest Intel® Architecture products and technologies.

Authored by core technology and product development team members, Intel Press books are among the earliest and most comprehensive means for developers to learn the inner workings of leading computer technologies.

Intel Press has gathered leaders in the field to open their notebooks and transmit their knowledge to the technology community. According to Intel Press, the only way to gain more direct access with industry experts would be in a face-to-face meeting—but even that couldn't match the comprehensive information that these technical books contain.

Intel Press is also sponsoring a well-stocked bookstore at the Intel Developer Forum Conference Spring '02. There, developers will find books recommended by IDF topic experts and organized by track from Intel Press and other major publishers. This will be a great opportunity to expand your technical reference library and prepare for continuing education after the conference.

New Book Titles

The new Intel Press books just released include:

- *The Virtual Interface Architecture*—Explains how a new interface solves the long-standing problem of efficiently interfacing general-purpose computers to high-speed switched networks
- *IXP1200 Programming: The Microengine Coding Guide for the Intel® IXP1200 Network Processor Family*—Learn ways to accomplish a complete range of typical tasks and understand how the hardware influences the programming environment
- *Video in the 21st Century*—Explains the state-of-the-art of digital video
- *The Software Optimization Cookbook: High-Performance Recipes for the Intel Architecture*—Explains how to optimize software for the Intel® Pentium® III and Pentium 4 processors
- *Building a Simple Network, Second Edition*—Learn how to connect a group of computers, printers, and communications devices in a small business or home office

All five books are detailed resources for those who create or deploy the leading edge of technology. Content ranges from fundamentals like schematics, source code, project files, and step-by-step guides to explanations of the strategy and concept behind the technology.

The Virtual Interface Architecture

Greg Regnier and Don Cameron's The Virtual Interface Architecture explains how a new interface solves the long-standing problem of efficiently interfacing general-purpose computers to high-speed switched networks, just as Virtual Memory allowed personal computers to break through physical memory limitations.

IXP1200 Programming: The Microengine Coding Guide for the Intel® IXP1200 Network Processor Family

In this concise, lively book on programming Intel IXP1200 processors, Erik J. Johnson and Aaron Kunze succinctly list ways to accomplish a complete range of typical tasks and tell you how the hardware influences the programming environment. As users progress from neophyte to veteran Intel IXP1200 processor programmers, insights in the book help with tuning Intel IXP1200 processor code for high performance.

Video in the 21st Century

Scott Janus's Video in the 21st Century describes the state-of-the-art of digital video in a manner understandable by a junior engineer, yet its information is still useful to a senior engineer. Video format standards, such as HDTV, DVD, and MPEG-2 are described, along with video display devices and protocols.

The Software Optimization Cookbook: High-Performance Recipes for the Intel® Architecture

Richard Gerber's book is about optimizing software for the Intel Pentium III and Pentium 4 processors. Topics include determining how fast the application currently runs, finding the areas of the program that need improvement, analyzing the program to determine what kinds of optimizations will help, general optimizations, optimizations that take advantage of specific processor features, and how to design for speed from the beginning.

Building a Simple Network, Second Edition

Ken Denniston's book provides a practical explanation of standard Ethernet networking products and tells how to connect a group of computers, printers, and communications devices in a small business or home office. With this book as your guide, you can install, set up, and use a network based on the Microsoft Windows* XP, Windows 2000, Windows Me, Windows 95, and Windows 98 operating systems.

Summary

With these five new titles, Intel Press expands its collection of first-to-market resources providing engineer-to-engineer technical information on the most recent Intel developments. Each book offers background on the underlying technology as well as practical directions and instructions for developing and optimizing computer systems for Intel architectures.

More Intel Press publications are under development, including comprehensive books on Intel® Itanium™ processor optimization and tuning, scientific and engineering computation on Intel Itanium processors, parallel programming with OpenMP, and network technologies.

Intel Press titles available now include:

- *The Virtual Interface Architecture* by Don Cameron and Greg Regnier
- *IXP1200 Programming* by Erik Johnson and Aaron Kunze
- *Video in the 21st Century* by Scott Janus
- *The Software Optimization Cookbook* by Richard Gerber
- *USB Design by Example, Second Edition* by John Hyde
- *Itanium™ Architecture for Software Developers* by WalterTriebel
- *InfiniBand* Architecture Development and Deployment* by William T. Futral
- *Programming Itanium™-based Systems* by Walter Triebel, Joe Bissell, and Rick Booth
- *Building the Power-Efficient PC* by Jerzy Kolinski, Barry Press, Ram Chary, and Andy Henroid
- *Peer-to-Peer Computing* by David Barkai
- *Building a Simple Network, Second Edition* by Ken Denniston
- *The Adaptive Enterprise* by Bruce Robertson and Valentin Sribar
- *Enriching the Value Chain* by Bruce Robertson and Valentin Sribar
- *Securing Business Information* by Dale Kutnick and F. Christian Byrnes

More Info

Intel Press books can be purchased directly from retail bookstores, including online retailers such as Fatbrain, Powell's, ShopIntel, Barnes & Noble, Borders, and Amazon.com.

To learn more about Intel Press and its publications, visit the Intel Press site.

Author Bio

Richard Bowles is the founding publisher and managing director of Intel Press. He has over 15 years of Intel tenure and has served in a wide variety of market, product, and sales development capacities.

Prior to founding Intel Press in 1999, Richard planned and directed the initial phase of Intel's market development in the technical workstation arena, and he founded the Intel Communications and Training Laboratory.

Richard has also held senior marketing and general management positions at Hewlett Packard, and Tektronix, and he has consulted to several start-up and Fortune 100 companies during more than 20 years in the electronics industry. Richard holds a Bachelor of Science degree in mechanical engineering from Stanford University, where he has completed post graduate studies in engineering management.

—End of Intel Developer Update Magazine Issue 29—